

# Contrail Release 5.0.3 Release Notes

Published  
2023-10-16

RELEASE

# Table of Contents

[Introduction](#)

[New and Changed Features](#)

[Supported Platforms Contrail 5.0](#)

[Known Behavior](#)

[Resolved Issues](#)

[Deprecated Items](#)

[Revision History](#)

# Introduction

Juniper Networks Contrail is an open, standards-based software solution that delivers network virtualization and service automation for federated cloud networks. It provides self-service provisioning, improves network troubleshooting and diagnostics, and enables service chaining for dynamic application environments across enterprise virtual private cloud (VPC), managed Infrastructure as a Service (IaaS), and Networks Functions Virtualization (NFV) use cases.

These release notes accompany Release 5.0.3 of Juniper Networks Contrail. They describe new features, limitations, and known problems.

These release notes are displayed on the Juniper Networks Contrail Documentation Web page at [https://www.juniper.net/documentation/en\\_US/contrail5.0/information-products /topic-collections/ release-notes/index.html](https://www.juniper.net/documentation/en_US/contrail5.0/information-products /topic-collections/ release-notes/index.html).

## New and Changed Features

### IN THIS SECTION

- [New and Changed Features in Contrail Release 5.0.3 | 1](#)
- [New and Changed Features in Contrail Release 5.0.2 | 2](#)
- [New and Changed Features in Contrail Release 5.0.1 | 4](#)
- [New and Changed Features in Contrail Release 5.0 | 12](#)

The features listed in this section are new or changed as of Contrail Release 5.0. A brief description of each new feature is included.

### New and Changed Features in Contrail Release 5.0.3

The features listed in this section are new as of Contrail Release 5.0.3.

## Support for Public Clouds

Contrail Release 5.0.3 supports Amazon Web Services (AWS) cloud with RHEL 7.5 instances.

## Contrail Enterprise Multicloud Deployment on Pre-Deployed RHEL 7.5 Compute Nodes

Contrail Release 5.0.3 supports Contrail Enterprise Multicloud (CEM) configuration on pre-deployed RHEL 7.5 compute nodes. For more information, see [Contrail Enterprise Multicloud Deployment using REST API](#).

## Documentation Update

Juniper Networks provides cumulative documentation for features supported in Contrail 5.0.X releases. Features supported only in Contrail releases 5.0.1, 5.0.2, and 5.0.3 are indicated in the respective topics in the documentation. You can also find feature support information in the release notes.

# New and Changed Features in Contrail Release 5.0.2

The features listed in this section are new as of Contrail Release 5.0.2.

## Support for Contrail Deployment on Microsoft Azure

Contrail Release 5.0.2 supports extending of on-premise Contrail capability on to Microsoft Azure public cloud. The multicloud gateway feature in Release 5.0.2 enables leveraging Contrail services and workloads to the public cloud seamlessly. The multicloud gateway provides seamless and secure connectivity between the on-premise system and Azure. Once secure connectivity is established, the on-premise workloads can be deployed on Azure. This workflow of spinning up Contrail SDN in a multicloud environment for Azure is automated. See [Deploying Contrail on Microsoft Azure](#) for information on deploying Contrail on Azure.

## Support for Mellanox Connectx-5 NIC on RHEL 7.5

Starting with Contrail Release 5.0.2, Contrail vRouter in DPDK mode supports the Mellanox Connectx-5 Network Interface Card (NIC) on the Red Hat Enterprise Linux 7.5 platforms. The NIC works in a no-offload mode in which all packets through the interface are transmitted to the vrouter-dpdk application and then sent to the respective virtual machines (VMs) or the host.

By default, DPDK mode uses the `uio_pci_generic` driver. In order to use a different driver, such as Mellanox, you must set the `ContrailDpdkDriver` heat template parameter in the **`contrail-services.yaml`** file.

```
parameter_defaults:  
  ContrailDpdkDriver: mlx
```

You must also, set the `driver` field to `mlx` in the NIC template configurations.

## Using Juju Charms to Deploy Contrail

You can deploy Contrail by using Juju Charms. Juju helps you deploy, configure, and efficiently manage applications on private clouds and public clouds. A Charm is a module containing a collection of scripts and metadata and is used with Juju to deploy Contrail. Juju Charms helps reduce the complexity of deploying Contrail by providing a simple way to deploy, configure, scale, and manage Contrail operations.

Starting with Release 5.0.2, Contrail supports the following charms:

- `contrail-agent`
- `contrail-analytics`
- `contrail-analyticsdb`
- `contrail-controller`
- `contrail-keystone-auth`
- `contrail-openstack`

For more information see, [Deploying Contrail by Using Juju Charms](#).

## Support for vSphere ESX Agent Manager

VMware provides a standard vCenter solution called vSphere ESX Agent Manager (EAM), that allows you to deploy, monitor, and manage ContrailVMs on ESXi hosts.

Starting in Contrail Release 5.0.2, the ContrailVM is deployed as an Agent VM that is monitored by EAM. With this integration, ContrailVMs are treated as more critical and privileged than other tenant VMs on the host.

For more information, see [vCenter Integration for Contrail Release 5.0.2](#).

## Support for Access Control Lists

Contrail release 5.0.2 supports conversion of security groups to access control lists (ACLs). These ACLs are applied to logical interfaces on QFX series devices that are connected to bare metal servers (BMS). In releases prior to Contrail Release 5.0.2, QFX supports only ingress filters on the Ethernet switching interfaces.

## Documentation Update

Juniper Networks provides cumulative documentation for features supported in Contrail Release 5.0, Contrail Release 5.0.1, and Contrail Release 5.0.2. Features supported only in Contrail releases 5.0.1 and 5.0.2 are indicated in the respective topics in the documentation. You can also find feature support information in the release notes.

## New and Changed Features in Contrail Release 5.0.1

The features listed in this section are new as of Contrail Release 5.0.1.

### Bare Metal Server Life Cycle Management

Starting with Contrail Release 5.0.1, you can use Bare Metal Server (BMS) Manager to manage the life cycle of bare metal servers (BMS). BMS Manager can install tenant-specific software images on BMS and attach them to the tenant network in a multi-tenant cloud. BMS Manager provides a single-click solution for the tenants to manage the bare metal servers in their network.

For more information, see [Understanding Bare Metal Server Managers](#).

### Contrail Support for Underlay Management

Contrail Release 5.0.1 supports underlay network management. The existing Contrail configuration node can provide intent driven automation capabilities on physical network elements such as ToR and EoR switches, Spines, SDN gateway, and VPN gateways in the data center. In addition, you can perform basic device management functions such as device discovery, device import, image upgrade, device underlay configuration, and topology discovery from the node.

For more information, see [Understanding Underlay Management](#).

### Support for vCenter-as-orchestrator Mode

Starting in Contrail Release 5.0.1, the only supported mode of vCenter integration with Contrail is the vCenter-as-orchestrator mode. With Release 5.0.1, Contrail no longer supports the vCenter-as-compute

mode, where the orchestrator is OpenStack, and the vCenter cluster acts as a nova-compute node to the OpenStack controller.

For more information, see [vCenter Integration for Contrail Release 5.0.1](#) and [Underlay Network Configuration for ContrailVM](#).

## Support for Advertising of Local AS to BGP Peers

Contrail Release 5.0.1 supports advertising of local autonomous system (AS) numbers to BGP peers. With the support for local AS, virtual network functions can advertise AS numbers of the received networks to Contrail, which results in the routing table getting updated and in turn enabling the sites to communicate. A new parameter `local-autonomous-system` has been added in the `BgpSessionAttributes` field in the Contrail configuration schema to support local AS. You can configure this parameter either from the Contrail web UI or by using the VNC API.

In earlier releases, Contrail supported advertising of only Contrail AS numbers, so the virtual machines could not advertise the AS numbers of received networks. This is because most BGP implementations do not advertise routes received from a given AS back to the same AS. With the support for local AS in Contrail, this check can be avoided, thus ensuring that VNF routes get correctly propagated by BGP.

## Support for Contrail Command UI

Contrail Release 5.0.1, supports the Contrail Command UI. The Contrail Command is an intuitive, wizard-based user interface (UI) which provides automated work flows such as the following:

- Contrail cluster deployment (Kolla-based OpenStack cluster)
- Automating the data center IP fabric
- Orchestrating virtual machines and bare metal servers

For more information, see [Configuring Contrail Command](#).

## Usage of Contrail Command UI

Use the Contrail command UI for only the following deployments and feature management:

- For Contrail DataCenter automation and fabric deployment and features, use the Contrail Command UI for both Contrail Release 5.0 and 5.0.1 features.
- For OpenStack-based Contrail Deployments, use the Contrail Command UI for both Contrail Release 5.0 and 5.0.1 features.
- For non-OpenStack based Contrail deployments, such as vCenter, Kubernetes, and OpenShift, use the old UI for features supported till Contrail Release 5.0.

- For non-OpenStack based Contrail deployments, such as vCenter, Kubernetes, and OpenShift, use VNC APIs for new features supported in Contrail Release 5.0.1.

### **Support for Configuring QFX10000 as a DataCenter Gateway**

Starting with Contrail Release 5.0.1, you can use a QFX10000 switch as a DataCenter Gateway (DC-GW). DC-GW is an overlay role that is assigned to a QFX10000 switch to:

- Extend private network
- Extend public routable network

For more information, see [Configuring QFX10000 as a DataCenter Gateway](#).

### **Support for Red Hat OpenShift Container Platform Version 3.9**

OpenShift Container Platform version 3.9 is supported on Contrail Release 5.0.1 and validated with RHEL 7.5.

For more information, see [Provisioning Red Hat OpenShift Container Platform Clusters Using Ansible Deployer](#).

### **Support for Contrail Cloud Release 13.0**

Contrail Release 5.0.1 supports Contrail Cloud Release 13.0. Contrail Cloud supports the OpenStack Queens Release and supports integration with Red Hat OpenStack Release 13. This OpenStack release provides support for containerized OpenStack services and Ansible-based RHOSP deployment.

For more information on Contrail Cloud Release 13.0, see [Contrail Cloud Platform Architecture](#) and [Contrail Cloud Deployment Guide](#).

### **Support for AppFormix Features**

Starting with Contrail Release 5.0.1, the following AppFormix features are integrated into the Contrail Command UI.

- Alarms
- Charts
- Contrail Control plane service monitoring
- Monitoring of vRouter stats

Prior to Contrail Release 5.0.1, Contrail Analytics and AppFormix were supported as separate packages and feature sets.



In Contrail 5.0.1, AppFormix provisioning is not supported using the UI, so you need to enable AppFormix in the Contrail Command UI.

To enable AppFormix in the Contrail Command UI:

1. Log in to the `contrail_command` Docker container.

```
docker exec -it contrail_command bash
```

2. Update the `feature-list.json` file and set the `appformix` key to `true`.

```
vi /usr/share/contrail/public/feature-list.json
```

3. Restart the `contrail_command` container.

```
docker restart contrail_command
```

4. Log in to Contrail Command UI and access the cluster endpoints.

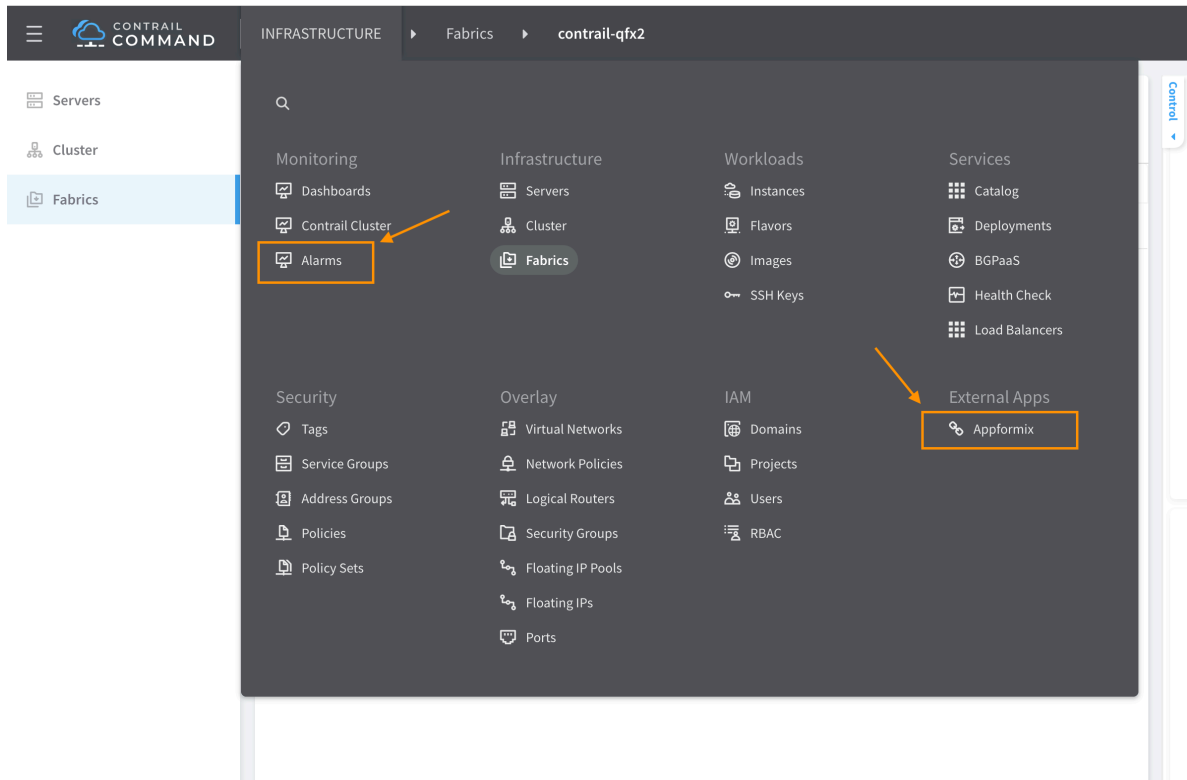
`http://<command-ui-ip>:9091/#/contrail_command/infrastructure/cluster-details/overview/advanced/endpoints`

5. Add a new entry with prefix `appformix` and update the `public_url` pointing to your Appformix UI IP:port.

You should now be able to access AppFormix features from within Contrail Command.

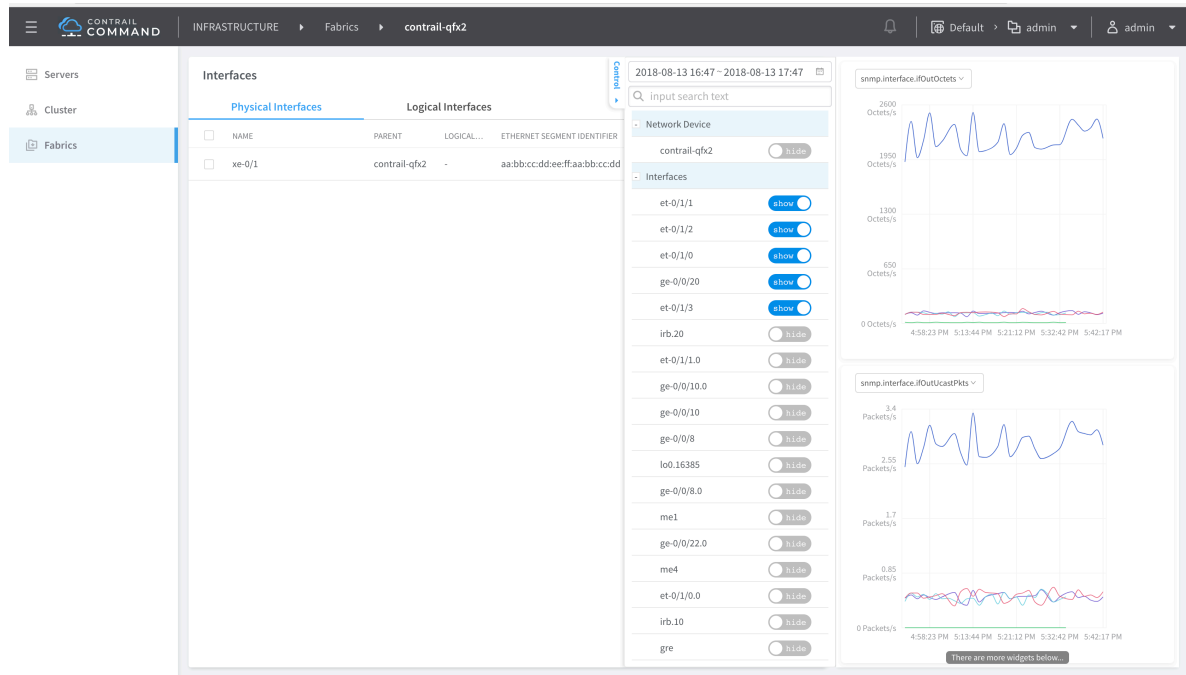
When enabled, you can see alarms under Monitoring and external link to Appformix UI under External Apps. See [Figure 1 on page 8](#).

Figure 1: AppFormix Features in Contrail Command UI



Additionally, network devices, instances, and host detail pages are augmented to show charts using metrics from AppFormix data. See [Figure 2 on page 9](#)

**Figure 2: Example of Network Devices in Contrail Command UI**



## Support for EVPN Route Type 5

Contrail Release 5.0.1 supports EVPN Route Type 5 messages as defined in the IETF specification *IP Prefix Advertisement in EVPN*. EVPN Route Type 5 is an extension of EVPN Route Type 2, which carries MAC addresses along with their associated IP addresses. EVPN Route Type 5 facilitates inter-subnet routing.

## Support For Encryption of Traffic Between vRouters

Contrail Release 5.0.1 supports encryption of data packets sent between Contrail vRouters. Contrail uses the IPsec Encapsulating Security Payload (ESP) protocol to provide authentication, integrity, and encryption to Layer 3 VPN (L3VPN) multi-tenant traffic egressing a vRouter. IPsec ESP ensures secure forwarding of tenant data in both private and public Contrail clouds.

You can enable encryption of multi-tenant traffic by enabling the `crypt_interface = crypt0` configuration option during provisioning. You can create IPsec tunnels between a set of vRouter compute nodes or across all vRouter compute nodes. In addition, you can specify a list of IP addresses that should send and receive encrypted multi-tenant traffic in the vRouter agent configuration file.

In releases prior to Contrail Release 5.0.1, L3VPN packets that egress a vRouter using GRE, VXLAN, or MPLS are not encrypted.

**NOTE:** This feature is supported only on On-premise data centers running Ubuntu release 16.04.

## Contrail Plugin For VMware vRealize Orchestrator

The Contrail plugin for VMware vRealize Orchestrator is available in Contrail Release 5.0.1. You can use the dedicated Contrail plugin to connect Contrail to VMware vRealize Orchestrator (vRO). vRO is used to automate the management processes in data centers. You can use the Contrail plugin to view the Contrail controller configurations in the vRO inventory. You can also use the plugin to modify configurations by using vRO workflows. You can deploy the Contrail plugin in any Java Virtual Machine (JVM) compatible language and load it on an active vRO instance. In Release 5.0.1, Contrail plugin for vRO does not support draft mode.

See [Integrating Contrail Release 5.0.1 with VMware vRealize Orchestrator](#)

## Support for Mellanox Connectx-5 NIC

Starting with Contrail Release 5.0.1, Contrail vRouter in DPDK mode supports the Mellanox Connectx-5 Network Interface Card (NIC). The NIC works in a no-offload mode in which all packets through the interface are transmitted to the vrouter-dpdk application and then sent to the respective virtual machines (VMs) or the host.

To deploy a node with the Mellanox Connectx-5 NIC, set "DPDK\_UIO\_DRIVER" : "mlx" for that node under vrouter in the **instances.yaml** or **host.yml** files. The vrouter-dpdk application cannot work if the "DPDK\_UIO\_DRIVER" is not set.

Example **instances.yaml** file in Ansible-based provisioning setups:

```
Bms1:
  provider: bms
  ip: 192.0.2.0
  roles:
    vrouter:
      AGENT_MODE: dpdk
      CPU_CORE_MASK: 0xff
      DPDK_UIO_DRIVER: mlx
      HUGE_PAGES: 32000
```

Example **host.yml** file in Helm-based provisioning setups:

```
AGENT_MODE: dpdk
CPU_CORE_MASK: 0xff
DPDK_UIO_DRIVER: mlx
HUGE_PAGES: 32000
```

## Support for Remote Compute

Contrail Release 5.0.1 supports remote compute, a method of managing a Contrail deployment across many small distributed data centers. Remote compute employs a subcluster that manages compute nodes at remote sites to receive configurations and exchange routes.

For more information, see [Remote Compute](#).

## Support for Red Hat OpenStack Platform Director 13

Contrail Release 5.0.1 supports integration with Red Hat OpenStack Platform Director 13.

[Table 1 on page 11](#) lists the the OpenStack releases and the corresponding operating systems and deployer versions supported by Contrail Release 5.0.1.

**Table 1: Supported Release Versions**

Contrail Release	Operating System	OpenStack	Deployer
Contrail 5.0.1	RHEL 7.5	Red Hat OpenStack Platform 13	RHOSP 13 director

For more information, see [Installing Contrail with Red Hat OpenStack Platform Director 13](#)

## Documentation Update

Juniper Networks provides cumulative documentation for features supported in Contrail Release 5.0 and Contrail Release 5.0.1. Features supported only in Contrail Release 5.0.1 are indicated in the respective topics in the documentation. You can also find feature support information in the release notes.

## New and Changed Features in Contrail Release 5.0

The features listed in this section are new as of Contrail Release 5.0.

### Ansible Scripts to Provision Contrail

Contrail Release 5.0 introduces microservices architecture. The `contrail-ansible-deployer` is a set of Ansible playbooks designed to deploy Contrail 5.x with microservices architecture on a CentOS-based system.

For more information, see [Overview of contrail-ansible-deployer for Installing Contrail with Microservices Architecture](#) and [Installing Contrail with OpenStack Ocata and Kolla Ansible](#).

### Contrail Microservices

With Contrail Release 4.0, Contrail started moving to an architecture of containers for major system components. Each container encapsulates the services needed for that container. The first phase of Contrail containers were characterized as fat containers, where multiple processes run within the container.

Starting with Contrail Release 5.0, more components are being containerized, and the fat containers are being decomposed into thin containers with microservices. The microservices are still encapsulated in their respective containers, however, only the essential functions relative to each container's functions are present as microservices. This enables a more agile system, avoiding monolithic containers.

Nothing is changing with regard to Contrail functionality, however, employing microservices provides a number of benefits, including the ability to deploy patches without updating the entire Contrail deployment, offering better ways to manage the lifecycles of containers, and improving user experiences with Contrail provisioning and upgrading. The microservices architecture enables provisioning with minimum information provided and enables every feature to be configurable. Utilizing microservices also simplifies application complexity by implementing small and independent processes.

For more information, see [Introduction to Contrail Microservices Architecture](#).

### Containerization of DPDK vRouter

Starting with Contrail Release 5.0, you can configure the Contrail DPDK vRouter to run in a Docker container. In earlier releases, DPDK vRouter runs on a compute host. The `contrail-vrouter-dpdk` binary file provides data plane functionality when Contrail vRouter is run in DPDK mode in a Contrail cluster.

For more information, see [Configuring Contrail DPDK vRouter to Run in a Docker Container](#).

## Distributed Source Network Address Translation (SNAT)

The distributed SNAT feature allows virtual machines to communicate with the IP fabric network using the existing forwarding infrastructure for compute node connectivity. This functionality is achieved through port address translation of virtual machine traffic using the IP address of the compute node as the public address.

The following distributed SNAT use case is supported:

- Virtual networks with distributed SNAT enabled can communicate with the IP fabric network. The session must be initiated from a virtual machine. Sessions initiated from the external network are not supported.

Distributed SNAT is supported only for TCP and UDP, and you can configure discrete port ranges for both protocols.

For more information, see [Source Network Address Translation \(SNAT\)](#).

## EVPN vRouter MultiHoming to Multiple ToRs

The Contrail control node can be in a situation in which it peers with a set of provider edge (PE) nodes that also contain a multihome CE device or top-of-rack (ToR) that is multihomed to the PE nodes, and the PEs are in all-active multihoming mode.

In this situation, any EVPN route originating through the multihome device is exported to Contrail with two paths—the multihome device path and the PE path. Those routes need to be load-balanced to prevent skewed traffic flow.

In previous releases, only the Contrail controller supports this scenario. In Contrail Release 5.0, this support has been added to the Contrail vRouter supporting both Layer 2 and Layer 3 traffic.

## Fat Flow Enhancements

The fat flow feature has been enhanced to support aggregation of multiple flows into a single flow by ignoring source and destination ports and/or IP addresses, or a combination of these. This extends the existing option of ignoring by either source or destination for a given protocol only.

Also added is support for fat-flow configuration at the VN level, extending the existing support at only the VMI level.

For more information, see [Understanding Flow Sampling](#).

## Implementing Kubernetes Network Policy with Contrail Firewall Policy

Contrail Release 5.0 supports implementing Kubernetes network policy in Contrail using the Contrail firewall security policy framework. While Kubernetes network policy can be implemented using other security objects in Contrail like security groups and Contrail network policies, the support of tags by Contrail firewall policy allows decoupling of routing from security policies and provides multi dimension segmentation and policy portability, while significantly enhancing user visibility and analytics functions.

For more information, see [Implementing Kubernetes Network Policy with Contrail Firewall Policy](#).

## Kubernetes Updates

Contrail Release 5.0 includes the following Kubernetes updates.

- The IP fabric forwarding feature enables reachability to public cloud services for Kubernetes pods. The IP fabric forwarding feature enables the overlay network to be a part of the underlay network or the IP fabric network, eliminating the need for encapsulating data packets between Kubernetes pods.
- The ip-fabric-snat feature enables service or ingress reachability from external clusters in isolated namespaces.
- Multiple Ingress Controllers can co-exist in Contrail. Since Contrail ensures the reachability between pods and services, any ingress controller can reach the endpoints or pods directly or through services.
- Contrail supports custom networks in namespace level. Starting with Contrail Release 5.0, custom networks are supported for ingress resources as well.
- Contrail network policy is created between the IP fabric network and pod-network to provide reachability between node and pods. So, any process in the node can reach the pods. Kubernetes Service Node-Port is also supported.

For more information on Kubernetes updates in Contrail Release 5.0, see [Kubernetes Updates](#).

## Routing Policies Enhanced for Interface Routes

Service interface and static routes are proliferating as leaked routes in the routing table of the SDN gateway. To reduce these leaked routes, routing policies have additional term match conditions under the protocol options to distinguish interface routes, service interface routes, and static routes from other VM routes. Also, a new action attribute, ASPATH (autonomous system path), is added that can be appended with a configurable AS list. All of the action attributes of Add/Set/Remove Community, SetLocal-Pref, and Set Med are supported with the new protocol match conditions and the new ASPATH list append action.



Users will be able to configure the new term match and action attributes as needed in the following cases:

- Setting LocalPref on service interface static routes when exporting to distinguish routes and take further action.
- Setting different LocalPref for all other reoriginated routes to distinguish routes and take further action.

Additionally, Contrail can set the LocalPref based on community onto imported routes, instead of the data center gateway, allowing direct access to VPN Internet-Shared from Contrail.

### **Service Instance Health Check Failure**

In Contrail Release 5.0, when one or more than one service instance (SI) in a service chain fails, reorigination of routes on the ingress and egress sides of the service chain is stopped. The routes automatically converge to a backup service chain that is part of another Contrail cluster. You can detect an SI failure by keeping track of corresponding connected routes of the service chain address.

For more information, see [Service Instance Health Checks](#).

### **Support for Load Balancing as a Service (LBaaS) in the Web UI**

For the LBaaS feature, load balancers using HAproxy can now be created, edited, or deleted using the Contrail Web UI.

For more information, see [Configuring Load Balancing as a Service in Contrail](#).

### **Support for Security Policies Draft Mode**

Starting with Contrail Release 5.0, you can define new security policies and review the policies before enforcing them. You can also edit existing policies and review the changes before updating them. You can define security policies in both global and project scopes.

For more information, see [Security Policies Draft Mode Overview](#).

### **Support for Virtual Network Route Tables in Contrail Introspect**

Starting with Contrail Release 5.0, virtual network route table entries per compute node can be viewed in Contrail Introspect.

## Support for a Flow-Hold Entries Counter in vRouter UVEs

Starting with Contrail Release 5.0 , a flow-hold entries counter is transmitted in vRouter User-Visible Entities (UVEs). The counter specifies the number of flows in **hold** state in the vRouter. vRouter uses the flow-hold count to check against a defined limit and when it reaches a defined limit, packets requiring new flows are dropped and new flows are not created till the flow-hold count goes below the defined limit.

Along with the flow-hold entries counter, the vRouter UVEs also provide information associated with a vRouter, such as:

- Virtual networks present on the vRouter
- Virtual machines spawned on the server of the vRouter
- Statistics of the traffic flowing through the vRouter

## Timestamp In UVE API Response

Starting with Contrail Release 5.0, a timestamp is added to the `/analytics/uve` UVE API response message.

## Timestamp In UVE Stream Response

Starting with Contrail Release 5.0, a timestamp is added to the `/analytics/uve-stream` UVE Stream API response message.

## Using Helm Charts to Provision Contrail

Starting with Contrail 5.0, Contrail Helm charts give you complete life cycle management of installation, update, and deletion of Contrail Docker-based containers in a microservices architecture.

Helm is the package manager for Kubernetes which is an open source software for managing containerized systems. The packaging format used by Helm is a chart, a collection of files that describe a related set of Kubernetes resources.

Many Contrail components have been broken out into manageable Helm charts, including the following specific features:

- Contrail service and IP address numbers are configurable by means of Helm charts.
- Ingress controllers can be implemented by means of Helm charts.

For more information, see:

- [Installing and Managing Contrail 5.0 Microservices Architecture Using Helm Charts](#)

- [Using Helm Charts to Provision Multinode Contrail OpenStack Ocata with High Availability](#)
- [Using Helm Charts to Provision All-in-One Contrail with OpenStack Ocata](#)
- [Accessing a Contrail OpenStack Helm Cluster](#)
- [Frequently Asked Questions About Contrail and Helm Charts](#)

## Contrail Plugin For VMware vRealize Orchestrator—Beta

The Contrail plugin for VMware vRealize Orchestrator is available as a Beta feature in Contrail Release 5.0. You can use the dedicated Contrail plugin to connect Contrail to VMware vRealize Orchestrator (vRO). vRO is used to automate the management processes in data centers. You can use the Contrail plugin to view the Contrail controller configurations in the vRO inventory. You can also use the plugin to modify configurations by using vRO workflows. You can deploy the Contrail plugin in any Java Virtual Machine (JVM) compatible language and load it on an active vRO instance.

See [Integrating Contrail with VMware vRealize Orchestrator](#).

## RBAC Support for Contrail Analytics API—Beta

Starting with Contrail Release 5.0, the Contrail Analytics API supports role-based access control (RBAC) as a Beta feature. Based on the user privileges, the logged-in user can access network monitoring information. Contrail Analytics API provides this information by mapping the user query and the UVE to the configuration objects on which RBAC rules are applied.

## Remote Compute—Beta

Remote compute is available as a Beta feature. The remote compute feature enables the deployment of Contrail in many small distributed data centers, up to hundreds or even thousands, for telecommunications point-of-presence (PoPs) or central offices (COs). Because each small datacenter has only a small number of computes running only a few applications, it is not cost-effective to deploy a full Contrail cluster of nodes of control, configuration, analytics, database, and the like, on dedicated servers in each distributed PoP. Additionally, manually managing hundreds or thousands of clusters is not feasible operationally.

Remote compute implements a subcluster that manages compute nodes at remote sites to receive configurations and exchange routes.

For more information, see [Remote Compute](#).

## Contrail Cloud

Contrail Cloud is *not* supported in Contrail Release 5.0.

## OpenShift Enterprise

OpenShift Enterprise is *not* supported in Contrail Release 5.0. OpenShift Origin 3.7 is supported.

# Supported Platforms Contrail 5.0

Table 2 on page 18 lists the orchestrator releases and the corresponding operating systems and kernel versions supported by Contrail Release 5.0.

**Table 2: Supported Platforms**

Contrail Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Release 5.0.3	Kubernetes 1.12	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> Docker version: 18.06.0-ce
	OpenShift 3.9	Ansible	<ul style="list-style-type: none"> <li>RHEL7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul>
	OpenStack Queens	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.6—Linux Kernel Version 3.10.0-957</li> </ul> Ansible version: 2.4.2 Docker version: 18.06.0-ce
	OpenStack Ocata	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.6—Linux Kernel Version 3.10.0-957</li> </ul> Ansible version: 2.4.2 Docker version: 18.06.0-ce

**Table 2: Supported Platforms (Continued)**

Contrail Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
		Helm	<ul style="list-style-type: none"> <li>• Ubuntu 16.04.3—Linux Kernel Version 4.4.0-112-generic</li> </ul> <p>Docker version: 17.03.2-ce</p> <p>Helm version: 2.7.2</p> <p>Kubernetes version: 1.9.3</p>
	Red Hat OpenStack Platform 13	RHOSP 13 director	<ul style="list-style-type: none"> <li>• RHEL7.6 - Linux Kernel Version 3.10.0-957</li> </ul>
	VMware vCenter 6.5	Ansible	<ul style="list-style-type: none"> <li>• ESX version 6.5</li> </ul> <p>CentOS VM version running vRouter: CentOS 7.5—Linux Kernel Version 3.10.0-862.9.1</p>
Contrail Release 5.0.2	Kubernetes 1.9.2	Ansible	<ul style="list-style-type: none"> <li>• CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> <p>Docker version: 18.06.0-ce</p>
	OpenShift 3.9	Ansible	<ul style="list-style-type: none"> <li>• CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> <p>Ansible version: 2.5.7</p> <p>Docker version: 1.13.1</p>
	OpenStack Queens	Ansible	<ul style="list-style-type: none"> <li>• CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> <p>Ansible version: 2.4.2</p> <p>Docker version: 18.06.0-ce</p>

**Table 2: Supported Platforms *(Continued)***

Contrail Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	OpenStack Ocata	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> Ansible version: 2.4.2 Docker version: 18.06.0-ce
		Helm	<ul style="list-style-type: none"> <li>Ubuntu 16.04.3—Linux Kernel Version 4.4.0-112-generic</li> </ul> Docker version: 17.03.2-ce Helm version: 2.7.2 Kubernetes version: 1.9.3
	Red Hat OpenStack Platform 13	RHOSP 13 director	<ul style="list-style-type: none"> <li>RHEL7.6 - Linux Kernel Version 3.10.0-957</li> </ul>
	VMware vCenter 6.5	Ansible	<ul style="list-style-type: none"> <li>ESX version 6.5</li> </ul> CentOS VM version running vRouter: CentOS 7.5—Linux Kernel Version 3.10.0-862.9.1
Contrail Release 5.0.1	Kubernetes 1.9.2	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> Docker version: 18.06.0-ce
	OpenShift 3.9	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> Ansible version: 2.5.7 Docker version: 1.13.1

**Table 2: Supported Platforms *(Continued)***

Contrail Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	OpenStack Queens	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> <p>Ansible version: 2.4.2</p> <p>Docker version: 18.06.0-ce</p>
	OpenStack Ocata	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul> <p>Ansible version: 2.4.2</p> <p>Docker version: 18.06.0-ce</p>
		Helm	<ul style="list-style-type: none"> <li>Ubuntu 16.04.3—Linux Kernel Version 4.4.0-112-generic</li> </ul> <p>Docker version: 17.03.2-ce</p> <p>Helm version: 2.7.2</p> <p>Kubernetes version: 1.9.3</p>
	Red Hat OpenStack Platform 13	RHOSP 13 director	<ul style="list-style-type: none"> <li>RHEL7.5—Linux Kernel Version 3.10.0-862.11.6</li> </ul>
	VMware vCenter 6.5	Ansible	<ul style="list-style-type: none"> <li>ESX version 6.5</li> </ul> <p>CentOS VM version running vRouter: CentOS 7.5—Linux Kernel Version 3.10.0-862.9.1</p>
Contrail Release 5.0	Kubernetes 1.9.2	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.4—Linux Kernel Version 3.10.0-693.21.1</li> </ul> <p>Docker version: 17.03.1-ce</p>

**Table 2: Supported Platforms *(Continued)***

Contrail Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	OpenShift 3.7	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.4—Linux Kernel Version 3.10.0-693.21.1</li> </ul> <p>Ansible version: 2.5.2</p> <p>Docker version: 1.13.1</p>
	OpenStack Ocata	Ansible	<ul style="list-style-type: none"> <li>CentOS 7.4—Linux Kernel Version 3.10.0-693.21.1</li> </ul> <p>Ansible version: 2.4.2</p> <p>Docker version: 17.03.1-ce</p>
		Helm	<ul style="list-style-type: none"> <li>Ubuntu 16.04.3—Linux Kernel Version 4.4.0-87</li> </ul> <p>Docker version: 1.13.1</p> <p>Helm version: 2.7.2</p> <p>Kubernetes version: 1.9.3</p>
	VMware vCenter 6.5	Ansible	<ul style="list-style-type: none"> <li>ESX version 6.5</li> </ul> <p>CentOS VM version running vRouter: CentOS 7.4—Linux Kernel Version 3.10.0-693.21.1</p>



# Known Behavior

## IN THIS SECTION

- [Known Behavior in Contrail Release 5.0.3 | 23](#)
- [Known Behavior in Contrail Release 5.0.2 | 28](#)
- [Known Behavior in Contrail Release 5.0.1 | 33](#)
- [Known Behavior in Contrail Release 5.0 | 36](#)

This section lists known limitations with this release.

## Known Behavior in Contrail Release 5.0.3

- CEM-3456 Keystone authentication works only with default domain. One known use case where non-default domain is used is Juju-Charm based integration with Bionic Queens OpenStack.
- 1803426 Between time intensive device manager operations (eg., device discovery, fabric role assignment, fabric delete, etc.) on the fabric, it is recommended to wait for sufficient time and make sure the previous operation is done.
- 1803253 BMS cannot be attached to Spine devices on the fabric.
- 1802683 VLAN tagged BMSes attached to same leaf node in a fabric, but under different virtual network must assume different VLANs.
- 1802425 If VxLAN mode is set to user mode, Contrail controller advertises routes without VxLAN encapsulation BGP community to the Route Reflector. This make the Route Reflector mark those routes as hidden. Hence it is recommended to use VxLAN mode auto.
- 1802420 MX gateway as a client to QFX is not supported. MX Gateway should also be configured as RR along with QFX.
- 1802060 When assigning roles to fabric devices, the Spine roles must be assigned before the Leaf roles.
- 1802000 After discovering a device, if the user want to delete the device, it can be done only after assigning a role to the device.

- 1801474 After provisioning, if the contrail-status shows that schema is in initializing state with connection to Cassandra down error, then restarting the schema service will recover from the error.

1801468 Security group configuration does not work on Leaf ports if QFX10K is configured as Leaf device.

- 1801401 While choosing security group names in fabric deployments, the name should be restricted to 26 characters as device manager prepends few fields before the name while applying on the devices.
- 1800207 While provisioning a k8s cluster, the k8s pods sometimes get stuck in ContainerCreate state. k8s cleanup with drain and delete of all nodes and re-join fixes the issue.
- 1799804 When extending the VN to Spine, the right way is to create a LR in Spine and attach the VN to the LR. Instead, if the user extends the VN to the Spine, the packets from the private network will be sent to the public network.
- 1799608 For a scaled fabric deployment, it is recommended to set max\_request\_size=10240000 in entrypt.sh in config container and restart the container. This is to allow for the huge size of message the DM pushes.
- 1799377 Adding new vnet to existing cloud on vnet under the same region is failing. The failure is that the gateways in the original vnet become unreachable. The public IPs of the original vnets are non-pingable and gateway ansible play fails with the unreachability error.

Workaround: Specify unique rules in the subnet of the new vnet being added.

- 1798914 SNAT feature does not work if VxLAN routing is enabled on the fabric.
- 1798684 While deleting a LR, the "remove" icon next to the LR listing does not work. The workaround is to use select the LR and then click the "remove" icon on the top of the page.
- 1797952 Before running play books for provisioning the user need to make sure docker-py package is not present on the system. If it is present, it can be removed by the following command: `pip uninstall docker-py docker`.
- 1797832 While bulk VM launch or VM migration on vcenter cluster, it is observed that sometimes vmware dvswitch sometimes ceases forwarding dhcp packets from guest vm to vrouter.

Workaround: Trigger some change to the portgroup.

For example: portgroup > edit settings > security > Promiscus mode > accept > OK. Then portgroup > edit settings > security > Promiscus mode > reject > OK. This trigger causes the dvswitch to start forwarding.

- 1794887 In Fabric ZTP scenario, dhcp\_server\_1 container in the all-in-one ZTP cluster will stuck in the restarting phase. This is because the dhcp subnet info is not present in the initial configuration

and this state will recover when the ZTP process is started from the UI which inturn applys the dhcp subnet configuration and restarts the container.

- 1794702 It is recommended to keep the number of community tags less than 50 in Network Policy.
- 1793811 The DHCP leases issued by the ZTP VM will be destroyed if the ZTP VM is brought down. To get away from this situation, it is recommended to have static IP config in DHCP bootstrap configuration in the ZTP VM.
- 1793269 k8s pod creation fails after config api restart To recover this, we will need to do the following:
  1. Login to the RabbitMQ docker on all the controllers, backup the directory “/var/lib/rabbitmq/mnesia/contrail@<Node\_Name>” and delete it.
  2. Once this folder is deleted in all the RabbitMQ dockers, Stop the rabbitmq docker on all the controllers.
  3. Start the rabbitmq docker in only one of the controller.
  4. Wait for about 10 secs, and restart the RabbitMQ docker on other two controller.
  5. To verify that the rabbitmq cluster is correct, execute the following command and verify that all three nodes are present in the field “running\_nodes” :

```
root@Config2:~/mnesia/contrail@Config2# rabbitmqctl cluster_status
Cluster status of node contrail@Config2
[{nodes,[{disc,[contrail@Config1,contrail@Config2,contrail@Config3]}]},
{running_nodes,[contrail@Config3,contrail@Config1,contrail@Config2]},
{cluster_name,<<"contrail@Config1.englab.juniper.net">>},
{partitions,[]},
{alarms,[{contrail@Config3,[]}, {contrail@Config1,[]}, {contrail@Config2,[]}]}]
```

- 1791244 In vcenter deployments, recommended procedure to bring a ESXi host out of standby mode:
  1. Resolve Issues from EAM - will exit the ESXi host from standby mode and power on ContrailVM.
  2. If ESXi host is powered on, wait for host to be back up and "Resolve Issues" from EAM post as applicable In the above procedure, if the EAM looses connection and the agency is deleted, apply the following workaround:

To resolve the issue as described in the bug, restart the ESX Agent Manager service on vcenter server. Administration > Deployment > System Configuration > Services > ESX Agent Manager.

Restart To tack and debug the EAM exception, SR is raised with VMWare - #18932826409

- 1787693 In case of BMS LCM the TFTP and PXE boot traffic should be passed to the openstack node. The default security group applied on the fabric devices blocks these and BMS LCM fails. To workaround this, add Ingress/0.0.0.0/ANY/ANY rule to the security group.
- 1803465 When BMS are connected across different subnets that are part of same VN, the required gateway configuration is not pushed to the Spine by the device manager. To get the traffic flow between the BMS in such scenario, the workaround is to configure static ARP on the BMS.
- 1803755 While creating ironic-provision, service address in the subnet should be pointing to OpenStack ironic node ip/kolla internal vip. Otherwise, ironic-provision vn will get configured with .2 as service address and it is seen that tftp requests are sent to TSN/CSN node which is .2.
- 1628326 DPDK-based compute node may not handle ARP packets correctly if only one CPU core is assigned to the vRouter in the test bed file.
- 1694343 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, you cannot set jumbo MTU size. Maximum MTU allowed: <=1500 . There is no workaround at present.
- 1714063 Analytics services does not support Keystone V3.
- 1716308 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- 1730021 In Contrail Security, non-admin users cannot access global scope objects when `global_access` is set to 0.

As a workaround, for non-admin users to access global scope objects, set `global_access` to 5.

- 1733027 In a multi-node analytics cluster, when one of analytics nodes shuts down or is rebooted, the contrail-status on the other analytics nodes displays the status of contrail-analytics-api and contrail-alarm-gen as initializing with the error message: Redis-UE:<ip-address of analytics node that was shutdown/rebooted>:6381[None] connection down. This does not impact the functioning of the analytics cluster.
- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.

- To install Pip, use the following command:

```
wget https://bootstrap.pypa.io/get-pip.py python get-pip.py
```

- To install docker-py, use the following command:

```
pip install docker-py
```

- 1749614 In a Helm-based provisioned cluster, VM launch fails if MariaDB replication is set to >1.
- 1754742 When you receive an error message during Kolla provisioning, rerunning the code will not work. In order for the provisioning to work, restart provisioning from scratch.
- 1755997 Kubernetes IP fabric feature is supported only on the pods and not the services.
- 1759428 Cannot provision an external MX series router by running the `provision_mx.py` script that contains the list of external routers.
- 1759576 Metadata SSL works only in HA deployment mode.
- 1759695 After deleting a vRouter chart with DPDK, the NICS do not rebind to the host in Helm.
- 1761137 High Availability provisioning of Kubernetes master is not supported.
- 1764610 Ansible-deployer: DPDK PPS is low.
- 1765281 On a DPDK compute `vif list -rate` core-dumps with traffic.
- 1766315 After provisioning Contrail by using a Helm-based provisioned cluster, restart nova-compute container.
- 1767472 SR-IOV with DPDK co-existence deployment is not supported using contrail-helm-deployer.
- 1782067 CFM: BGP peering configuration from an existing device is not removed while unassigning the physical device.
- 1784750 Contrail RBAC (API and analytics) policies are not part of Contrail Helm charts unlike in OpenStack components. You cannot override global/domain/project RBAC policies while deploying Helm charts.
- 1786348 Contrail-Command: Token expires immediately after clicking **Create Instance** on the UI. This causes the instance to be present in the UI, but entire data does not get pushed in the backend.

- 1786527 In Helm-based provisioning when the vRouter gets provisioned, the vRouter gets disconnected from other vRouters residing on different subnets. As a workaround, move the host routes used to reach different subnets to vhost0 interface and restart the system.
- 1786560 In Kubernetes single yaml file-based provisioning, when the vRouter gets provisioned, the vRouter gets disconnected from other vRouters residing on different subnets. As a workaround, move the host routes used to reach different subnets to vhost0 interface and restart the system.
- 1787303 Contrail Release 5.0.1 uses DPDK version 17.11.03, DPDK vRouter using uio driver as vfio-pci, typically for Intel Fortville NICs, and needs firmware upgrade to the following version or later (if it is certified by users).

driver: i40e

version: 2.1.14-k

firmware-version: 6.01 0x80003493 255.65535.255

## Known Behavior in Contrail Release 5.0.2

- 1803426 Between time intensive device manager operations (eg., device discovery, fabric role assignment, fabric delete, etc.) on the fabric, it is recommended to wait for sufficient time and make sure the previous operation is done.
- 1803253 BMS cannot be attached to Spine devices on the fabric.
- 1802683 VLAN tagged BMSes attached to same leaf node in a fabric, but under different virtual network must assume different VLANs.
- 1802425 If VxLAN mode is set to user mode, Contrail controller advertises routes without VxLAN encapsulation BGP community to the Route Reflector. This make the Route Reflector mark those routes as hidden. Hence it is recommended to use VxLAN mode auto.
- 1802420 MX gateway as a client to QFX is not supported. MX Gateway should also be configured as RR along with QFX.
- 1802060 When assigning roles to fabric devices, the Spine roles must be assigned before the Leaf roles.
- 1802000 After discovering a device, if the user want to delete the device, it can be done only after assigning a role to the device.
- 1801474 After provisioning, if the contrail-status shows that schema is in initializing state with connection to Cassandra down error, then restarting the schema service will recover from the error.

1801468 Security group configuration does not work on Leaf ports if QFX10K is configured as Leaf device.

- 1801401 While choosing security group names in fabric deployments, the name should be restricted to 26 characters as device manager prepends few fields before the name while applying on the devices.
- 1800207 While provisioning a k8s cluster, the k8s pods sometimes get stuck in ContainerCreate state. k8s cleanup with drain and delete of all nodes and re-join fixes the issue.
- 1799804 When extending the VN to Spine, the right way is to create a LR in Spine and attach the VN to the LR. Instead, if the user extends the VN to the Spine, the packets from the private network will be sent to the public network.
- 1799608 For a scaled fabric deployment, it is recommended to set max\_request\_size=10240000 in entrypt.sh in config container and restart the container. This is to allow for the huge size of message the DM pushes.
- 1799377 Adding new vnet to existing cloud on vnet under the same region is failing. The failure is that the gateways in the original vnet become unreachable. The public IPs of the original vnets are non-pingable and gateway ansible play fails with the unreachability error.

Workaround: Specify unique rules in the subnet of the new vnet being added.

- 1798914 SNAT feature does not work if VxLAN routing is enabled on the fabric.
- 1798684 While deleting a LR, the "remove" icon next to the LR listing does not work. The workaround is to use select the LR and then click the "remove" icon on the top of the page.
- 1797952 Before running play books for provisioning the user need to make sure docker-py package is not present on the system. If it is present, it can be removed by the following command: `pip uninstall docker-py docker`.
- 1797832 While bulk VM launch or VM migration on vcenter cluster, it is observed that sometimes vmware dvswitch sometimes ceases forwarding dhcp packets from guest vm to vrouter.

Workaround: Trigger some change to the portgroup.

For example: portgroup > edit settings > security > Promiscus mode > accept > OK. Then portgroup > edit settings > security > Promiscus mode > reject > OK. This trigger causes the dvswitch to start forwarding.

- 1794887 In Fabric ZTP scenario, dhcp\_server\_1 container in the all-in-one ZTP cluster will stuck in the restarting phase. This is because the dhcp subnet info is not present in the initial configuration and this state will recover when the ZTP process is started from the UI which inturn applys the dhcp subnet configuration and restarts the container.
- 1794702 It is recommended to keep the number of community tags less than 50 in Network Policy.

- 1793811 The DHCP leases issued by the ZTP VM will be destroyed if the ZTP VM is brought down. To get away from this situation, it is recommended to have static IP config in DHCP bootstrap configuration in the ZTP VM.
- 1793269 k8s pod creation fails after config api restart To recover this, we will need to do the following:
  1. Login to the RabbitMQ docker on all the controllers, backup the directory “/var/lib/rabbitmq/mnesia/contrail@<Node\_Name>” and delete it.
  2. Once this folder is deleted in all the RabbitMQ dockers, Stop the rabbitmq docker on all the controllers.
  3. Start the rabbitmq docker in only one of the controller.
  4. Wait for about 10 secs, and restart the RabbitMQ docker on other two controller.
  5. To verify that the rabbitmq cluster is correct, execute the following command and verify that all three nodes are present in the field “running\_nodes” :

```
root@Config2:~/mnesia/contrail@Config2# rabbitmqctl cluster_status
Cluster status of node contrail@Config2
[{nodes,[{disc,[contrail@Config1,contrail@Config2,contrail@Config3]}]},
{running_nodes,[contrail@Config3,contrail@Config1,contrail@Config2]},
{cluster_name,<<"contrail@Config1.englab.juniper.net">>},
{partitions,[[]]},
{alarms,[{contrail@Config3,[[]]}, {contrail@Config1,[[]]}, {contrail@Config2,[[]]}]}
```

- 1791244 In vcenter deployments, recommended procedure to bring a ESXi host out of standby mode:
  1. Resolve Issues from EAM - will exit the ESXi host from standby mode and power on ContrailVM.
  2. If ESXi host is powered on, wait for host to be back up and "Resolve Issues" from EAM post as applicable In the above procedure, if the EAM loses connection and the agency is deleted, apply the following workaround:
 

To resolve the issue as described in the bug, restart the ESX Agent Manager service on vcenter server. Administration > Deployment > System Configuration > Services > ESX Agent Manager.

Restart To track and debug the EAM exception, SR is raised with VMWare - #18932826409
- 1787693 In case of BMS LCM the TFTP and PXE boot traffic should be passed to the openstack node. The default security group applied on the fabric devices blocks these and BMS LCM fails. To workaround this, add Ingress/0.0.0.0/ANY/ANY rule to the security group.



- 1803465 When BMS are connected across different subnets that are part of same VN, the required gateway configuration is not pushed to the Spine by the device manager. To get the traffic flow between the BMS in such scenario, the workaround is to configure static ARP on the BMS.
- 1803755 While creating ironic-provision, service address in the subnet should be pointing to OpenStack ironic node ip/kolla internal vip. Otherwise, ironic-provision vn will get configured with .2 as service address and it seen that tftp requests are sent to TSN/CSN node which is .2.
- 1628326 DPDK-based compute node may not handle ARP packets correctly if only one CPU core is assigned to the vRouter in the test bed file.
- 1694343 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, you cannot set jumbo MTU size. Maximum MTU allowed:  $\leq 1500$ . There is no workaround at present.
- 1714063 Analytics services does not support Keystone V3.
- 1716308 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- 1730021 In Contrail Security, non-admin users cannot access global scope objects when `global_access` is set to 0.

As a workaround, for non-admin users to access global scope objects, set `global_access` to 5.

- 1733027 In a multi-node analytics cluster, when one of analytics nodes shuts down or is rebooted, the contrail-status on the other analytics nodes displays the status of contrail-analytics-api and contrail-alarm-gen as initializing with the error message: Redis-UVE:<ip-address of analytics node that was shutdown/rebooted>:6381[None] connection down. This does not impact the functioning of the analytics cluster.
- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.
  - To install Pip, use the following command:

```
wget https://bootstrap.pypa.io/get-pip.py python get-pip.py
```

- To install docker-py, use the following command:

```
pip install docker-py
```

- 1749614 In a Helm-based provisioned cluster, VM launch fails if MariaDB replication is set to >1.
- 1754742 When you receive an error message during Kolla provisioning, rerunning the code will not work. In order for the provisioning to work, restart provisioning from scratch.
- 1755997 Kubernetes IP fabric feature is supported only on the pods and not the services.
- 1759428 Cannot provision an external MX series router by running the provision\_mx.py script that contains the list of external routers.
- 1759576 Metadata SSL works only in HA deployment mode.
- 1759695 After deleting a vRouter chart with DPDK, the NICS do not rebind to the host in Helm.
- 1761137 High Availability provisioning of Kubernetes master is not supported.
- 1764610 Ansible-deployer: DPDK PPS is low.
- 1765281 On a DPDK compute vif list -rate core-dumps with traffic.
- 1766315 After provisioning Contrail by using a Helm-based provisioned cluster, restart nova-compute container.
- 1767472 SR-IOV with DPDK co-existence deployment is not supported using contrail-helm-deployer.
- 1782067 CFM: BGP peering configuration from an existing device is not removed while unassigning the physical device.
- 1784750 Contrail RBAC (API and analytics) policies are not part of Contrail Helm charts unlike in OpenStack components. You cannot override global/domain/project RBAC policies while deploying Helm charts.
- 1786348 Contrail-Command: Token expires immediately after clicking **Create Instance** on the UI. This causes the instance to be present in the UI, but entire data does not get pushed in the backend.
- 1786527 In Helm-based provisioning when the vRouter gets provisioned, the vRouter gets disconnected from other vRouters residing on different subnets. As a workaround, move the host routes used to reach different subnets to vhost0 interface and restart the system.
- 1786560 In Kubernetes single yaml file-based provisioning, when the vRouter gets provisioned, the vRouter gets disconnected from other vRouters residing on different subnets. As a workaround, move the host routes used to reach different subnets to vhost0 interface and restart the system.

- 1787303 Contrail Release 5.0.1 uses DPDK version 17.11.03, DPDK vRouter using uio driver as vfio-pci, typically for Intel Fortville NICs, and needs firmware upgrade to the following version or later (if it is certified by users).

driver: i40e

version: 2.1.14-k

firmware-version: 6.01 0x80003493 255.65535.255

## Known Behavior in Contrail Release 5.0.1

- 1628326 DPDK-based compute node may not handle ARP packets correctly if only one CPU core is assigned to the vRouter in the test bed file.
- 1694343 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, you cannot set jumbo MTU size. Maximum MTU allowed:  $\leq 1500$ . There is no workaround at present.
- 1714063 Analytics services does not support Keystone V3.
- 1716308 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- 1730021 In Contrail Security, non-admin users cannot access global scope objects when `global_access` is set to `0`.

As a workaround, for non-admin users to access global scope objects, set `global_access` to `5`.

- 1733027 In a multi-node analytics cluster, when one of analytics nodes shuts down or is rebooted, the contrail-status on the other analytics nodes displays the status of contrail-analytics-api and contrail-alarm-gen as initializing with the error message: Redis-UVE:<ip-address of analytics node that was shutdown/rebooted>:6381[None] connection down. This does not impact the functioning of the analytics cluster.

- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.
- To install Pip, use the following command:

```
wget https://bootstrap.pypa.io/get-pip.py python get-pip.py
```

- To install docker-py, use the following command:

```
pip install docker-py
```

- 1749614 In a Helm-based provisioned cluster, VM launch fails if MariaDB replication is set to >1.
- 1754742 When you receive an error message during Kolla provisioning, rerunning the code will not work. In order for the provisioning to work, restart provisioning from scratch.
- 1755997 Kubernetes IP fabric feature is supported only on the pods and not the services.
- 1759428 Cannot provision an external MX series router by running the provision\_mx.py script that contains the list of external routers.
- 1759576 Metadata SSL works only in HA deployment mode.
- 1759695 After deleting a vRouter chart with DPDK, the NICS do not rebind to the host in Helm.
- 1761137 High Availability provisioning of Kubernetes master is not supported.
- 1764610 Ansible-deployer: DPDK PPS is low.
- 1764739 Docker restart hangs indefinitely and Docker daemon stops running post the hang.
- 1765281 On a DPDK compute vif list -rate core-dumps with traffic.
- 1766035 On a Kubernetes cluster, a controller node reboot fails to re-establish the BGP XMPP connection with compute nodes. As a workaround, flush the iptables on the Kubernetes master.
- 1766315 After provisioning Contrail by using a Helm-based provisioned cluster, restart nova-compute container.
- 1767472 SR-IOV with DPDK co-existence deployment is not supported using contrail-helm-deployer.
- 1771444 CFM: The MTU on devices between the IRONIC node and BMS must be manually configured accounting for the VxLAN encapsulation that carries traffic between them.
- 1777562 CFM: CSN HA is not supported

- 1778279 CFM: ESI Multihoming with multiple interfaces connected to the same ESI (same TOR) is not supported.
- 1781305 OpenShift deployer code does not have the option to choose different interface names for different physical nodes.
- 1782067 CFM: BGP peering configuration from an existing device is not removed while unassigning the physical device.
- 1784085 OpenShift deployments do not install NTP on the nodes. After provisioning, you must install and start the ntpd on all Openshift nodes manually.
- 1784750 Contrail RBAC (API and analytics) policies are not part of Contrail Helm charts unlike in OpenStack components. You cannot override global/domain/project RBAC policies while deploying Helm charts.
- 1785448 You cannot configure routing tables on the Contrail Command UI. Routing tables must be configured on the older Contrail UI at **Networking > Routing**.
- 1785925 CFM: Upon deletion of device images, a few chunks remain in the image repository.
- 1786102 CFM: In Contrail Release 5.0.1 release, interface channelization of devices in greenfield workflow is not supported. In brownfield, interface channelization is supported as long as it is configured manually before onboarding devices.
- 1786348 Contrail-Command: Token expires immediately after clicking **Create Instance** on the UI. This causes the instance to be present in the UI, but entire data does not get pushed in the backend.
- 1786354 CFM: Unassigning and reassigning the device to fabric doesn't reconfigure the QFX with Contrail configuration.
- 1786536 CFM: For auto configure of devices to be successful, while adding the devices, both physical and overlay roles must be specified.
- 1786487 Dynamic mirroring without Juniper header does not work for VxLAN encapsulation.
- 1786511 rabbitmq cluster formation fails during Helm sanity job. As a workaround, update the rabbitmq section of the **contrail-analytics/values.yaml** and **contrail-controller/values.yaml** files, to include the IP addresses of the controller nodes. For example, change `endpoints.rabbitmq.hosts.default: rabbitmq` to `endpoints.rabbitmq.hosts.default: "controller_node1_IP,controller_node2_IP,controller_node3_IP"`.
- 1786527 In Helm-based provisioning when the vRouter gets provisioned, the vRouter gets disconnected from other vRouters residing on different subnets. As a workaround, move the host routes used to reach different subnets to vhost0 interface and restart the system.

- 1786560 In Kubernetes single yaml file-based provisioning, when the vRouter gets provisioned, the vRouter gets disconnected from other vRouters residing on different subnets. As a workaround, move the host routes used to reach different subnets to vhost0 interface and restart the system.
- 1786836 CFM: Logical Router pushed to spine devices during device role assignment will not allow VxLAN routing to be enabled. To prevent this situation, configure VxLAN routing on spines before assigning the spine role on a physical router
- 1786855 CFM: BMS - LAG can be configured on a BMS only if life cycle management is executed through CFM.
- 1786856 Fabric device discovery fails intermittently in the Contrail HA topology. As a workaround, restart Config\_api\_1 on all HA nodes and run device discovery again.
- 1787303 Contrail Release 5.0.1 uses DPDK version 17.11.03, DPDK vRouter using uio driver as vfio-pci, typically for Intel Fortville NICs, and needs firmware upgrade to the following version or later (if it is certified by users).

driver: i40e

version: 2.1.14-k

firmware-version: 6.01 0x80003493 255.65535.255

- 1789768 Instances.yml in the contrail\_command container is populated with the default password instead of the one provided by the user.

## Known Behavior in Contrail Release 5.0

- 1694343 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, you cannot set jumbo MTU size. Maximum MTU allowed: <=1500 . There is no workaround at present.
- 1716308 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- 1730021 In Contrail Security, non-admin users cannot access global scope objects when `global_access` is set to 0.

As a workaround, for non-admin users to access global scope objects, set `global_access` to 5.

- 1733027 In a multi-node analytics cluster, when one of analytics nodes shuts down or is rebooted, the contrail-status on the other analytics nodes displays the status of contrail-analytics-api and contrail-alarm-gen as initializing with the error message: Redis-UE:<ip-address of analytics node that was shutdown/rebooted>:6381[None] connection down. This does not impact the functioning of the analytics cluster.
- 1735590 In Kubernetes and OpenShift based deployments when we crate SNAT router and extend cluster-network to that SNAT router host is losing all connectivity.

As a workaround, if you want to use the SNAT feature in Contrail, disassociate the ip-fabric-cluster-network-default policy and delete it.

- 1749614 In a Helm-based provisioned cluster, VM launch fails if MariaDB replication is set to >1.
- 1754742 When you receive an error message during Kolla provisioning, rerunning the code will not work. In order for the provisioning to work, restart provisioning from scratch.
- 1755997 Kubernetes IP fabric feature is supported only on the pods and not the services.
- 1759576 Metadata SSL works only in HA deployment mode.
- 1759695 After deleting a vRouter chart with DPDK, the NICS do not rebind to the host in Helm.
- 1761137 High Availability provisioning of Kubernetes master is not supported.
- 1764739 Docker restart hangs indefinitely and Docker daemon stops running post the hang.
- 1764925 RabbitMQ clustering fails on certain nodes in a setup. As a workaround, restart the container that is not in the cluster.
- 1765277 When a snapshot of an active VM fails, shutdown the VM before generating the snapshot.
- 1765281 On a DPDK compute vif list -rate core-dumps with traffic.
- 1765487 A false alarm for config service is generated when config and configdb services are installed on different nodes. Ignore the false alarm.
- 1766035 On a Kubernetes cluster, a controller node reboot fails to re-establish the BGP XMPP connection with compute nodes. As a workaround, flush the iptables on the Kubernetes master.
- 1766315 After provisioning Contrail by using a Helm-based provisioned cluster, restart nova-compute container.
- 1766371 OpenShift use cases work in non-HA environments only.

- 1767094 Kube DNS fails to come online come up in a multi-interface setup.
- 1767466 In Contrail 5.0 release, with Contrail Helm charts, Kubernetes ingress Web UI URL does not work when Web UI is started with secure (TLS) option.
- 1767470 SR-IOV installation is not supported with contrail-helm-deployer.
- 1767472 SR-IOV with DPDK co-existence deployment is not supported using contrail-helm-deployer.
- 1759428 Cannot provision an external MX series router by running the provision\_mx.py script that contains the list of external routers.

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues in Contrail Release 5.0.2 | 38](#)
- [Resolved Issues in Contrail Release 5.0.1 | 38](#)
- [Resolved Issues in Contrail Release 5.0 | 39](#)

This section lists limitations that are resolved with this release.

### Resolved Issues in Contrail Release 5.0.2

You can research limitations that are resolved with Contrail Release 5.0.2 in Launchpad at:

<https://launchpad.net/juniperopenstack/+milestone/r5.0.2>.

### Resolved Issues in Contrail Release 5.0.1

You can research limitations that are resolved with Contrail Release 5.0.1 in Launchpad at:

<https://launchpad.net/juniperopenstack/+milestone/r5.0.1>.



## Resolved Issues in Contrail Release 5.0

You can research limitations that are resolved with Contrail Release 5.0 in Launchpad at:

<https://launchpad.net/juniperopenstack/r5.0/r5.0.0>.

## Deprecated Items

The following features have been deprecated in Contrail Release 5.0.

- Contrail Server Manager
- Virtual Extensible LAN (VXLAN) and the Open vSwitch Database (OVSDB) management protocol based ToR and ToR services node (TSN) solution for bare metal-workload support is deprecated. Customers are recommended to use EVPN-VXLAN based Contrail Service Node solution which is more scalable.

## Revision History

March, 2019—Revision 1—Contrail 5.0.3.

---

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2023 Juniper Networks, Inc. All rights reserved.